EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	0	mapping same ((hardware adj accelerator) and (((executable binary) near4 (portion module kernal section block function))(processing near2 element near3 array)))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/03/17 15:45
L2	24	(mapping matrix) and ((hardware adj accelerator) and (((executable binary) near4 (portion module kernal section block function))(processing near2 element near3 array)))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/03/17 15:41
L3	17	2 and array	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/03/17 15:43
L4	5	3 and ((hardware near4 accelerator) and array and matrix)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/03/17 15:44
L5	1	4 and @rlad<"20000803"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/03/17 15:49
L6	0	(mapping and ((hardware adj accelerator) and (((executable binary) near4 (portion module kernal section block function))(processing near2 element near3 array)))).clm.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/03/17 15:45
L7	0	(variant and ((hardware adj accelerator) and (((executable binary) near4 (portion module kernal section block function))(processing near2 element near3 array)))).clm.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/03/17 15:46
L8	1	(((hardware adj accelerator) and ((processing near2 element near3 array)))).clm.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/03/17 15:47

EAST Search History

L9	23	(((executable binary) near4 (portion module kernal section block function)) and (matrix same (mapping referenc$4))).clm.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/03/17 15:47
L10	0	9 and ((processing near2 element near3 array) FPGA PGA PE) and accelerator	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/03/17 15:48
L11	33	(((processing near2 element near3 array) FPGA PGA PE) and accelerator).clm.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/03/17 15:48
L12	4	(9 or 11) and @rlad<"20000803"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/03/17 15:49

P☉RTAL
USPTO

"hardware accelerators" <and> partitioning <and> mapping <   **SEARCH**

THE ACM DIGITAL LIBRARY

┇ Feedback Report a problem Satisfaction survey

Terms used **hardware accelerators** and **partitioning** and **mapping** and **matrix**

Found **198,206** of **198,617**

Sort results by	relevance ▼	● Save results to a Binder	Try an Advanced Search
Display results	expanded form ▼	⸮ Search Tips	Try this search in The ACM Guide
		☐ Open results in a new window	

Results 1 - 20 of 200   . Result page: **1** 2 3 4 5 6 7 8 9 10   next
Best 200 shown   Relevance scale ☐▭▬▮■

1 Compilation: Scalable subgraph mapping for acyclic computation accelerators   ■
Nathan Clark, Amir Hormati, Scott Mahlke, Sami Yehia
October 2006 **Proceedings of the 2006 international conference on Compilers, architecture and synthesis for embedded systems CASES '06**
Publisher: ACM Press
Full text available: 📄 pdf(906.08 KB)   Additional Information: full citation, abstract, references, index terms

Computer architects are constantly faced with the need to improve performance and increase the efficiency of computation in their designs. To this end, it is increasingly common to see acyclic com-putation accelerators appear in embedded processor designs. One major problem with adding accelerators to a design is that it is difficult to generate high-quality code utilizing them. Hand-written assembly code is typical, and if compiler support does exist, it is implemented using only greedy algorit ...

Keywords: compilation, embedded processors

2 Pre-computed radiance transfer: theory and practice: Precomputed radiance transfer: ■
theory and practice
Jan Kautz, Peter-Pike Sloan, Jaakko Lehtinen
July 2005 **ACM SIGGRAPH 2005 Courses SIGGRAPH '05**
Publisher: ACM Press
Full text available: 📄 pdf(8.77 MB)   Additional Information: full citation, abstract, references

Interactive rendering of realistic objects under general lighting models poses three principal challenges. Handling complex light transport phenomena like shadows, inter-reflections, caustics and sub-surface scattering is difficult to do in real time. Integrating these effects over large area light sources compounds the difficulty, and finally real objects have complex spatially-varying BRDF's. Precomputed Radiance Transfer (PRT) encapsulates a family of techniques that partially addresses these ...

3 Technical papers: Data management and query---Hypergraph partitioning for   ■
automatic memory hierarchy management
Sriram Krishnamoorthy, Umit Catalyurek, Jarek Nieplocha, Atanas Rountev, P. Sadayappan
November 2006 **Proceedings of the 2006 ACM/IEEE conference on Supercomputing SC '06**
Publisher: ACM Press

Full text available: ⬛ pdf(351.32 KB)   Additional Information: full citation, abstract, references
⬛ html(2.18 KB)

In this paper, we present a mechanism for automatic management of the memory hierarchy, including secondary storage, in the context of a global address space parallel programming framework. The programmer specifies the parallelism and locality in the computation. The scheduling of the computation into stages, together with the movement of the associated data between secondary storage and global memory, and between global memory and local memory, is automatically managed. A novel formulation of h ...

4 Technical papers: Molecular dynamics---Preliminary investigation of advanced electrostatics in molecular dynamics on reconfigurable computers ⬛

Ronald Scrofano, Viktor K. Prasanna
November 2006 **Proceedings of the 2006 ACM/IEEE conference on Supercomputing SC '06**
Publisher: ACM Press
Full text available: ⬛ pdf(233.59 KB)   Additional Information: full citation, abstract, references
⬛ html(2.36 KB)

Scientific computing is marked by applications with very high performance demands. As technology has improved, reconfigurable hardware has become a viable platform to provide application acceleration, even for floating-point-intensive scientific applications. Now, reconfigurable computers---computers with general purpose microprocessors, reconfigurable hardware, memory, and high performance interconnect---are emerging as platforms that allow complete applications to be partitioned into parts tha ...

Keywords: FPGA, electrostatics, molecular dynamics, reconfigurable

5 The elements of nature: interactive and realistic techniques ⬛

Oliver Deusen, David S. Ebert, Ron Fedkiw, F. Kenton Musgrave, Przemyslaw Prusinkiewicz, Doug Roble, Jos Stam, Jerry Tessendorf
August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**
Publisher: ACM Press
Full text available: ⬛ pdf(17.65 MB)   Additional Information: full citation, abstract

This updated course on simulating natural phenomena will cover the latest research and production techniques for simulating most of the elements of nature. The presenters will provide movie production, interactive simulation, and research perspectives on the difficult task of photorealistic modeling, rendering, and animation of natural phenomena. The course offers a nice balance of the latest interactive graphics hardware-based simulation techniques and the latest physics-based simulation techni ...

6 Constructing and exploiting linear schedules with prescribed parallelism ⬛

Alain Darte, Robert Schreiber, B. Ramakrishna Rau, Frédéric Vivien
January 2002 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 7 Issue 1
Publisher: ACM Press
Full text available: ⬛ pdf(159.04 KB)   Additional Information: full citation, abstract, references, citings, index terms

We present two new results of importance in code generation for and synthesis of synchronously scheduled parallel processor arrays and multicluster VLIWs. The first is a new practical method for constructing a linear schedule for the iterations of a loop nest that schedules precisely one iteration per cycle on each of a prescribed set of processors. While this problem goes back to the era in which systolic computation was in vogue, it has defied practical solution until now. We provide a closed ...

Keywords: Linear schedule, multicluster VLIW, systolic array

7 Surface modeling and parameterization with manifolds: Surface modeling and parameterization with manifolds: Siggraph 2006 course notes
Author presenation videos are available from the citation page
Cindy Grimm, Denis Zorin
July 2006 **ACM SIGGRAPH 2006 Courses SIGGRAPH '06**
Publisher: ACM Press
Full text available: pdf(17.85 MB)
mov(251.00   Additional Information: full citation, abstract, references
bytes)

Many diverse applications in different areas of computer graphics, including geometric modeling, rendering and animation, require dealing with sets which cannot be easily represented with a single function on a simple domain in a Euclidean space: Examples include surfaces of nontrivial topology, environment maps, reflection/transmission functions, light fields, configuration spaces of animation skeletons, and others. In most cases these objects are described as collections of functions defined o ...

8 Manifolds and modeling: Surface modeling and parameterization with manifolds
Cindy Grimm, Denis Zorin
July 2005 **ACM SIGGRAPH 2005 Courses SIGGRAPH '05**
Publisher: ACM Press
Full text available: pdf(6.69 MB)   Additional Information: full citation, references

9 Active pages: a computation model for intelligent memory
Mark Oskin, Frederic T. Chong, Timothy Sherwood
April 1998 **ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture ISCA '98**, Volume 26 Issue 3
Publisher: IEEE Computer Society, ACM Press
Full text available: pdf(1.58 MB)   Additional Information: full citation, abstract, references, citings, index
Publisher Site   terms

Microprocessors and memory systems suffer from a growing gap in performance. We introduce *Active Pages,* a computation model which addresses this gap by shifting data-intensive computations to the memory system. An Active Page consists of a page of data and a set of associated functions which can operate upon that data. We describe an implementation of Active Pages on RADram (Reconfigurable Architecture DRAM), a memory system based upon the integration of DRAM and reconfigurable logic. Res ...

10 Design space exploration using arithmetic-level hardware--software cosimulation for configurable multiprocessor platforms
Jingzhao Ou, Viktor K. Prasanna
May 2006 **ACM Transactions on Embedded Computing Systems (TECS)**, Volume 5 Issue 2
Publisher: ACM Press
Full text available: pdf(814.20 KB)   Additional Information: full citation, abstract, references, index terms

Configurable multiprocessor platforms consist of multiple soft processors configured on FPGA devices. They have become an attractive choice for implementing many computing applications. In addition to the various ways of distributing software execution among the multiple soft processors, the application designer can customize soft processors and the

connections between them in order to improve the performance of the applications running on the multiprocessor platform. State-of-the-art design too ...

Keywords: FPGA, cosimulation, design space exploration, processor

11 Feature Selection for Unsupervised and Supervised Inference: The Emergence of Sparsity in a Weight-Based Approach
Lior Wolf, Amnon Shashua
December 2005 **The Journal of Machine Learning Research**, Volume 6
Publisher: MIT Press
Full text available: pdf(462.32 KB)   Additional Information: full citation, abstract

The problem of selecting a subset of relevant features in a potentially overwhelming quantity of data is classic and found in many branches of science. Examples in computer vision, text processing and more recently bio-informatics are abundant. In text classification tasks, for example, it is not uncommon to have 10^4 to 10^7 features of the size of the vocabulary containing word frequency counts, with the expectation that only a small fraction of them are relevant. Typical e ...

12 An adaptive and dynamic dimensionality reduction method for high-dimensional indexing
Heng Tao Shen, Xiaofang Zhou, Aoying Zhou
January 2007 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 16 Issue 2
Publisher: Springer-Verlag New York, Inc.
Full text available: pdf(570.24 KB)   Additional Information: full citation, abstract

The notorious "dimensionality curse" is a well-known phenomenon for any multi-dimensional indexes attempting to scale up to high dimensions. One well-known approach to overcome degradation in performance with respect to increasing dimensions is to reduce the dimensionality of the original dataset before constructing the index. However, identifying the correlation among the dimensions and effectively reducing them are challenging tasks. In this paper, we present an adaptive *Multi* ...

Keywords: *Correlated clustering, Dimensionality reduction, High-dimensional indexing, Projection, Subspace*

13 Area and delay estimation for FPGA implementation of coarse-grained reconfigurable architectures
Leipo Yan, Thambipillai Srikanthan, Niu Gang
June 2006 **ACM SIGPLAN Notices , Proceedings of the 2006 ACM SIGPLAN/SIGBED conference on Language, compilers and tool support for embedded systems LCTES '06**, Volume 41 Issue 7
Publisher: ACM Press
Full text available: pdf(144.89 KB)   Additional Information: full citation, abstract, references, index terms

Reconfigurable architecture is one solution to the increasing computational requirement that often cannot be met by the low-end embedded processors. Compiling applications to such architectures involves hardware/software partitioning. To partition the applications, a set of parameters, such as the hardware execution time and hardware area consumption, is required for each application block. Quick derivation of the parameters for all the blocks is essential. Previous research has shown that the c ...

Keywords: CGRA, VLIW, area estimation, delay estimation, hardware/software partitioning

14 Technical papers: Biology---Locality and parallelism optimization for dynamic programming algorithm in bioinformatics
Guangming Tan, Shengzhong Feng, Ninghui Sun
November 2006 **Proceedings of the 2006 ACM/IEEE conference on Supercomputing SC '06**
Publisher: ACM Press
Full text available: pdf(298.74 KB)   Additional Information: full citation, abstract, references
html(2.31 KB)

Dynamic programming has been one of the most efficient approaches to sequence analysis and structure prediction in biology. However, their performance is limited due to the drastic increase in both the number of biological data and variety of the computer architectures. With regard to such predicament, this paper creates excellent algorithms aimed at addressing the challenges of improving memory efficiency and network latency tolerance for nonserial polyadic dynamic programming where the depende ...

Keywords: cache-oblivious, dynamic programming, locality, parallelism, tiling

15 Compilers: Applications of storage mapping optimization to register promotion
Patrick Carribault, Albert Cohen
June 2004 **Proceedings of the 18th annual international conference on Supercomputing ICS '04**
Publisher: ACM Press
Full text available: pdf(268.41 KB)   Additional Information: full citation, abstract, references, index terms

Storage mapping optimization is a flexible approach to folding array dimensions in numerical codes. It is designed to reduce the memory footprint after a wide spectrum of loop transformations, whether based on uniform dependence vectors or more expressive polyhedral abstractions. Conversely, few loop transformations have been proposed to facilitate register promotion, namely loop fusion, unroll-and-jam or tiling. Building on array data-flow analysis and expansion, we extend storage mapping optim ...

Keywords: array contraction, array folding, blocking, itanium, pattern matching, register promotion, scheduling, string matching, tiling

16 Concepts and effectiveness of the cover-coefficient-based clustering methodology for text databases
Fazli Can, Esen A. Ozkarahan
December 1990 **ACM Transactions on Database Systems (TODS)**, Volume 15 Issue 4
Publisher: ACM Press

Full text available: pdf(2.74 MB)   Additional Information: full citation, abstract, references, citings, index terms, review

A new algorithm for document clustering is introduced. The base concept of the algorithm, the cover coefficient (CC) concept, provides a means of estimating the number of clusters within a document database and related indexing and clustering analytically. The CC concept is used also to identify the cluster seeds and to form clusters with these seeds. It is shown that the complexity of the clustering process is very low. The retrieval experiments show that the information-retrieval effectiv ...

Keywords: cluster validity, clustering-indexing relationships, cover coefficient, decoupling coefficient, document retrieval, retrieval effectiveness

17 <u>Reconfigurable system: Energy/power estimation of regular processor arrays</u>

Steven Derrien, Sanjay Rajopadhye

October 2002 **Proceedings of the 15th international symposium on System Synthesis ISSS '02**

Publisher: ACM Press

Full text available: pdf(909.01 KB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>index terms</u>

We propose a high-level analytical model for estimating the energy and/or power dissipation in VLSI processor (systolic) array implementations of loop programs, particularly for implementations on FPGA based CO-processors. We focus on the respective impact of the array design parameters on the overall off-chip i/o traffic and the number and sizes of the local memories in the array. The model is validated experimentally and shows good results (12.7% RMS error in the predictions).

Keywords: design space exploration, power estimation, processor array partitioning, programmable logic

18 <u>Multitasking on reconfigurable architectures: microarchitecture support and dynamic scheduling</u>

Juanjo Noguera, Rosa M. Badia

May 2004 **ACM Transactions on Embedded Computing Systems (TECS)**, Volume 3 Issue 2

Publisher: ACM Press

Full text available: pdf(1.18 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>

Dynamic scheduling for system-on-chip (SoC) platforms has become an important field of research due to the emerging range of applications with dynamic behavior (e.g., MPEG-4). Dynamically reconfigurable architectures are an interesting solution for this type of applications. Scheduling for dynamically reconfigurable architectures might be classified in two major broad categories: (1) static scheduling techniques or (2) use of an operating system (OS) for reconfigurable computing. However, resear ...

Keywords: Adaptable architectures and microarchitectures, dynamic scheduling, runtime support for dynamic reconfiguration

19 <u>Formal methods: Formal specification and verification of data separation in a separation kernel for an embedded system</u>

Constance L. Heitmeyer, Myla Archer, Elizabeth I. Leonard, John McLean

October 2006 **Proceedings of the 13th ACM conference on Computer and communications security CCS '06**

Publisher: ACM Press

Full text available: pdf(285.74 KB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>index terms</u>

Although many algorithms, hardware designs, and security protocols have been formally verified, formal verification of the security of software is still rare. This is due in large part to the large size of software, which results in huge costs for verification. This paper describes a novel and practical approach to formally establishing the security of code. The approach begins with a well-defined set of security properties and, based on the properties, constructs a compact security model contai ...

Keywords: code verification, formal model, formal specification, separation kernel, theorem proving

20 A formal method for hardware IP design and integration under I/O and timing constraints

Philippe Coussy, Emmanuel Casseau, Pierre Bomel, Adel Baganne, Eric Martin
February 2006 **ACM Transactions on Embedded Computing Systems (TECS)**, Volume 5 Issue 1
Publisher: ACM Press
Full text available: pdf(2.18 MB)   Additional Information: full citation, abstract, references, index terms

IP integration, which is one of the most important SoC design steps, requires taking into account communication and timing constraints. In that context, design and reuse can be improved using IP cores described at a high abstraction level. In this paper, we present an IP design approach that relies on three main phases: (1) constraint modeling, (2) IP constraint analysis steps for feasibility checking, and (3) synthesis. We propose a set of techniques dedicated to the digital signal processing d ...

Keywords: IP design and integration, SoC, communication interface unit, constrained synthesis, digital signal processing and multimedia applications

P⬤RTAL
USPTO

THE ACM DIGITAL LIBRARY

❧ Feedback Report a problem Satisfaction survey

Terms used   **Found**
FPGA processing element array and accelerator and code partitioning and mapping and matrix   **160,761 of 198,617**

Sort results by   [relevance ▼]   ❧ Save results to a Binder   Try an Advanced Search
Display results   [expanded form ▼]   ⓘ Search Tips   Try this search in The ACM Guide
                                       ☐ Open results in a new window

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10   next
Best 200 shown   Relevance scale ☐ ▭ ▬ ▰ ■

1 **A decade of reconfigurable computing: a visionary retrospective**   ■
    R. Hartenstein
    March 2001 **Proceedings of the conference on Design, automation and test in Europe DATE '01**
    Publisher: IEEE Press
    Full text available: 📄 pdf(768.00 KB)   Additional Information: full citation, references, citings, index terms

2 **Coarse grain reconfigurable architecture (embedded tutorial)**   ■
    Reiner Hartenstein
    January 2001 **Proceedings of the 2001 conference on Asia South Pacific design automation ASP-DAC '01**
    Publisher: ACM Press

    Full text available: 📄 pdf(167.05 KB)   Additional Information: full citation, abstract, references, citings, index terms

    The paper gives a brief survey over a decade of R&D on coarse grain reconfigurable hardware and related compilation techniques and points out its significance to the emerging discipline of reconfigurable computing.

3 **Active pages: a computation model for intelligent memory**   ■
    Mark Oskin, Frederic T. Chong, Timothy Sherwood
    April 1998 **ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture ISCA '98**, Volume 26 Issue 3
    Publisher: IEEE Computer Society, ACM Press
    Full text available: 📄 pdf(1.58 MB) 📄 Publisher Site   Additional Information: full citation, abstract, references, citings, index terms

    Microprocessors and memory systems suffer from a growing gap in performance. We introduce *Active Pages,* a computation model which addresses this gap by shifting data-intensive computations to the memory system. An Active Page consists of a page of data and a set of associated functions which can operate upon that data. We describe an implementation of Active Pages on RADram (Reconfigurable Architecture DRAM), a memory system based upon the integration of DRAM and reconfigurable logic. Res ...

4 A compiler approach to fast hardware design space exploration in FPGA-based systems

Byoungro So, Mary W. Hall, Pedro C. Diniz

May 2002 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 2002 Conference on Programming language design and implementation PLDI '02**, Volume 37 Issue 5

Publisher: ACM Press

Full text available: 🔂 pdf(359.71 KB)   Additional Information: full citation, abstract, references, citings, index terms

The current practice of mapping computations to custom hardware implementations requires programmers to assume the role of hardware designers. In tuning the performance of their hardware implementation, designers manually apply loop transformations such as loop unrolling. designers manually apply loop transformations. For example, loop unrolling is used to expose instruction-level parallelism at the expense of more hardware resources for concurrent operator evaluation. Because unrolling also inc ...

Keywords: data dependence analysis, design space exploration, loop transformations, reuse analysis

5 Design space exploration using arithmetic-level hardware--software cosimulation for configurable multiprocessor platforms

Jingzhao Ou, Viktor K. Prasanna

May 2006 **ACM Transactions on Embedded Computing Systems (TECS)**, Volume 5 Issue 2

Publisher: ACM Press

Full text available: 🔂 pdf(814.20 KB)   Additional Information: full citation, abstract, references, index terms

Configurable multiprocessor platforms consist of multiple soft processors configured on FPGA devices. They have become an attractive choice for implementing many computing applications. In addition to the various ways of distributing software execution among the multiple soft processors, the application designer can customize soft processors and the connections between them in order to improve the performance of the applications running on the multiprocessor platform. State-of-the-art design too ...

Keywords: FPGA, cosimulation, design space exploration, processor

6 Computation techniques for FPGAs: An FPGA-based VLIW processor with custom hardware execution

Alex K. Jones, Raymond Hoare, Dara Kusic, Joshua Fazekas, John Foster

February 2005 **Proceedings of the 2005 ACM/SIGDA 13th international symposium on Field-programmable gate arrays FPGA '05**

Publisher: ACM Press

Full text available: 🔂 pdf(220.52 KB)   Additional Information: full citation, abstract, references, citings, index terms

The capability and heterogeneity of new FPGA (Field Programmable Gate Array) devices continues to increase with each new line of devices. Efficiently programming these devices is increasing in difficulty. However, FPGAs continue to be utilized for algorithms traditionally targeted to embedded DSP microprocessors such as signal and image processing applications.This paper presents an architecture that combines VLIW (Very Large Instruction Word) processing with the capability to introduce applicat ...

Keywords: NIOS, VLIW, compiler, kernels, parallelism, synthesis

7 <u>High Performance Linear Algebra Operations on Reconfigurable Systems</u>

Ling Zhuo, Viktor K. Prasanna

November 2005 **Proceedings of the 2005 ACM/IEEE conference on Supercomputing SC '05**

Publisher: IEEE Computer Society

Full text available: <u>pdf(460.47 KB)</u>   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>index terms</u>

> Field-Programmable Gate Arrays (FPGAs) have become an attractive option for scientific computing. Several vendors have developed high performance reconfigurable systems which employ FPGAs for application acceleration. In this paper, we propose a BLAS (Basic Linear Algebra Subprograms) library for state-of-the-art reconfigurable systems. We study three data-intensive operations: dot product, matrix-vector multiply and dense matrix multiply. The first two operations are I/O bound, and our designs ...

8 <u>Evaluation of the streams-C C-to-FPGA compiler: an applications perspective</u>

Jan Frigo, Maya Gokhale, Dominique Lavenier

February 2001 **Proceedings of the 2001 ACM/SIGDA ninth international symposium on Field programmable gate arrays FPGA '01**

Publisher: ACM Press

Full text available: <u>pdf(180.81 KB)</u>   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>

> The Streams-C compiler ([5]) synthesizes hardware circuits for reconfigurable FPGA-based computers from parallel C programs. The Streams-C language consists of a small number of libraries and intrinsic functions added to a synthesizable subset of C, and supports a communicating process programming model. The processes may be either software or hardware processes, and the compiler manages communication among the processes transparently to the programmer. For the hardware processes, the compi ...

Keywords: FPGA, FPGA design tools, configurable computing, hardware-software co-design, high-level synthesis, silicon compiler

9 <u>Novel FPGA applications: Hyper customized processors for bio-sequence database scanning on FPGAs</u>

Tim Oliver, Bertil Schmidt, Douglas Maskell

February 2005 **Proceedings of the 2005 ACM/SIGDA 13th international symposium on Field-programmable gate arrays FPGA '05**

Publisher: ACM Press

Full text available: <u>pdf(304.99 KB)</u>   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>index terms</u>

> Protein sequences with unknown functionality are often compared to a set of known sequences to detect functional similarities. Efficient dynamic-programming algorithms exist for solving this problem, however current solutions still require significant scan times. These scan time requirements are likely to become even more severe due to exponential database growth. In this paper we present a new approach to bio-sequence database scanning using re-configurable FPGA-based hardware platforms to gain ...

Keywords: FPGA, Smith-Waterman, bio-informatics, dynamic re-configuration

10 <u>Real-time shading</u>

Marc Olano, Kurt Akeley, John C. Hart, Wolfgang Heidrich, Michael McCool, Jason L. Mitchell, Randi Rost

August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**

Publisher: ACM Press

Full text available: 🗎 pdf(7.39 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>

Real-time procedural shading was once seen as a distant dream. When the first version of this course was offered four years ago, real-time shading was possible, but only with one-of-a-kind hardware or by combining the effects of tens to hundreds of rendering passes. Today, almost every new computer comes with graphics hardware capable of interactively executing shaders of thousands to tens of thousands of instructions. This course has been redesigned to address today's real-time shading capabili ...

11 <u>Static array storage optimization in MATLAB</u> ■

Pramod G. Joisha, Prithviraj Banerjee
May 2003 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 2003 conference on Programming language design and implementation PLDI '03**, Volume 38 Issue 5
Publisher: ACM Press

Full text available: 🗎 pdf(287.36 KB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>

Static array storage optimization in MATLAB.

12 <u>Accurate Area and Delay Estimators for FPGAs</u> ■

A. Nayak, M. Haldar, A. Choudhary, P. Banerjee
March 2002 **Proceedings of the conference on Design, automation and test in Europe DATE '02**
Publisher: IEEE Computer Society

Full text available: 🗎 pdf(167.38 KB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>citings</u>

We present an area and delay estimator in the context of a compilerthat takes in high level signal and image processing applicationsdescribed in MATLAB and performs automatic design spaceexploration to synthesize hardware for a Field Programmable GateArray (FPGA) which meets the user area and frequency specifications.We present an area estimator which is used to estimatethe maximum number of Configurable Logic Blocks (CLBs) consumedby the hardware synthesized for the Xilinx XC4010 fromthe input ...

13 <u>Compilers: Applications of storage mapping optimization to register promotion</u> ■

Patrick Carribault, Albert Cohen
June 2004 **Proceedings of the 18th annual international conference on Supercomputing ICS '04**
Publisher: ACM Press

Full text available: 🗎 pdf(268.41 KB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>index terms</u>

Storage mapping optimization is a flexible approach to folding array dimensions in numerical codes. It is designed to reduce the memory footprint after a wide spectrum of loop transformations, whether based on uniform dependence vectors or more expressive polyhedral abstractions. Conversely, few loop transformations have been proposed to facilitate register promotion, namely loop fusion, unroll-and-jam or tiling. Building on array data-flow analysis and expansion, we extend storage mapping optim ...

Keywords: array contraction, array folding, blocking, itanium, pattern matching, register promotion, scheduling, string matching, tiling

14 <u>Factoring large numbers with programmable hardware</u> ■

Hea Joung Kim, William H. Mangione-Smith
February 2000 **Proceedings of the 2000 ACM/SIGDA eighth international symposium on Field programmable gate arrays FPGA '00**
Publisher: ACM Press

Full text available: 📄 pdf(760.09 KB)   Additional Information: full citation, abstract, references, citings, index terms

Most advanced forms of security for electronic transactions rely on the public-key cryptosystems developed by Rivest, Shamir and Adleman. Unfortunately, these systems are only secure while it remains difficult to factor large integers. The fastest published algorithms for factoring large numbers have a common sieving step. These sieves collect numbers that are completely factored by a set of prime numbers that are known in advance. Furthermore, the time required to execute these sieves curr ...

Keywords: configurable computing technologies, number factoring algorithms, public-key cryptosystems, sieving, special-purpose factoring hardware

15 Performance Analysis in CoDe-X Partitioning for Structural Programmable Accelerators
Reiner W. Hartenstein, Juergen Becker
March 1997 **Proceedings of the 5th International Workshop on Hardware/Software Co-Design CODES '97**
Publisher: IEEE Computer Society
Full text available: 📄 pdf(923.52 KB)
📄 Publisher Site
Additional Information: full citation, abstract

The paper presents the performance analysis process within the parallelizing compilation environment CoDe-X for simultaneous programming of Xputer-based accelerators and their host. The paper introduces briefly its hardware/software co-design strategies at two levels of partitioning. CoDe-X performs both, at first level a profiling-driven host/accelerator partitioning for performance optimization, and at second level a resource-driven sequential/structural partitioning of the accelerator source ...

Keywords: design space exploration, performance estimation, structural programmable co-processors

16 Combinational logic synthesis for LUT based field programmable gate arrays
Jason Cong, Yuzheng Ding
April 1996 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 1 Issue 2
Publisher: ACM Press
Full text available: 📄 pdf(628.91 KB)   Additional Information: full citation, abstract, references, citings, index terms, review

The increasing popularity of the field programmable gate-array (FPGA) technology has generated a great deal of interest in the algorithmic study and tool development for FPGA-specific design automation problems. The most widely used FPGAs are LUT based FPGAs, in which the basic logic element is a K-input one-output lookup-table (LUT) that can implement any Boolean function of up to K variables. This unique feature of the LUT has brought new challenges to lo ...

Keywords: FPGA, area minimization, computer-aided design of VLSI, decomposition, delay minimization, delay modeling, logic optimization, power minimization, programmable logic, routing, simplification, synthesis, system design, technology mapping

17 Optimizing compiler for shared-memory multiple SIMD architecture
Weihua Zhang, Xinglong Qian, Ye Wang, Binyu Zang, Chuanqi Zhu
June 2006 **ACM SIGPLAN Notices , Proceedings of the 2006 ACM SIGPLAN/SIGBED**

conference on Language, compilers and tool support for embedded systems LCTES '06, Volume 41 Issue 7
Publisher: ACM Press
Full text available: pdf(375.75 KB)   Additional Information: full citation, abstract, references, index terms

With the rapid growth of multimedia and game, these applications put more and more pressure on the processing ability of modern processors. Multiple SIMD architecture is widely used in multimedia processing field as a multimedia accelerator. With the consideration of power consumption and chip size, shared memory multiple SIMD architecture is mainly used in embedded SOCs. In order to further fit mobile environment, there is the constraint of limited register number as well. Although shared memory ...

Keywords: locality, multiple SIMD, optimization, replication, shared memory

18 Automatic compilation to a coarse-grained reconfigurable system-opn-chip

Girish Venkataramani, Walid Najjar, Fadi Kurdahi, Nader Bagherzadeh, Wim Bohm, Jeff Hammes
November 2003 **ACM Transactions on Embedded Computing Systems (TECS)**, Volume 2 Issue 4
Publisher: ACM Press

Full text available: pdf(687.52 KB)   Additional Information: full citation, abstract, references, citings, index terms

The rapid growth of device densities on silicon has made it feasible to deploy reconfigurable hardware as a highly parallel computing platform. However, one of the obstacles to the wider acceptance of this technology is its programmability. The application needs to be programmed in hardware description languages or an assembly equivalent, whereas most application programmers are used to the algorithmic programming paradigm. SA-C has been proposed as an expression-oriented language designed to im ...

Keywords: Reconfigurable computing, SIMD, compilers

19 Storage assignment during high-level synthesis for configurable architectures

Wenrui Gong, Gang Wang, R. Kastner
May 2005 **Proceedings of the 2005 IEEE/ACM International conference on Computer-aided design ICCAD '05**
Publisher: IEEE Computer Society
Full text available: pdf(202.77 KB)   Additional Information: full citation, abstract

Modern, high performance configurable architectures integrate on-chip, distributed block RAM modules to provide ample data storage. Synthesizing applications to these complex systems requires an effective and efficient approach to conduct data partitioning and storage assignment. In this paper, we present a data and iteration space partitioning solution that focuses on minimizing remote memory accesses or, equivalently, maximizing the local computation. Using the same code but different data par ...

20 Wireless communication and networking: Interconnection framework for high-throughput, flexible LDPC decoders

Federico Quaglio, Fabrizio Vacca, Cristiano Castellano, Alberto Tarable, Guido Masera
March 2006 **Proceedings of the conference on Design, automation and test in Europe: Designers' forum DATE '06**
Publisher: European Design and Automation Association
Full text available: pdf(196.54 KB)   Additional Information: full citation, abstract, references

This paper presents a possible interconnection structure suitable for being used in a flexible LDPC decoder. The main feature of the proposed approach is the possibility of

implementing parallel or semi-parallel decoders with a reduced communication complexity. To the best of our knowledge this is the first work detailing the implementation of a fully flexible LDPC decoder, able to support any type of code. To prove the effectiveness of this approach, a complete decoder has been implemented on a ...

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

Terms of Usage   Privacy Policy   Code of Ethics   Contact Us

Useful downloads: Adobe Acrobat   QuickTime   Windows Media Player   Real Player